(12) UK Patent Application (19) GB (11) 2 357 656 (13) A

(43) Date of A Publication 27.06.2001

(21) Application No 9916680.3

(22) Date of Filing 16.07.1999

(51) INT CL[7]
H04Q 11/00

(52) UK CL (Edition S)
H4K KTKX
H4B BN

(56) Documents Cited
None

(58) Field of Search
UK CL (Edition S) H4B BN , H4K KTKX
INT CL[7] H04Q 11/00
Online:WPI,EPODOC,JAPIO,INSPEC

(71) Applicant(s)
Opera Systems Limited
(Incorporated in the United Kingdom)
124 High Street, Midsomer Norton, BATH, BA3 2DA, United Kingdom

(72) Inventor(s)
Trevor James Hall
Robert Walter Alister Scarr
William Alden Crossland
Timothy David Wilkinson
Michael Anthony Hands
Richard George Hoptroff
Martin Hamilton George

(74) Agent and/or Address for Service
Opera Systems Limited
P.O.Box 2563, BATH, BA1 3YR, United Kingdom

(54) Abstract Title
Optical packet router

(57) A packet routing architecture consists of an input packet handler, a 'core' switch and an output packet handler. The input packet handler translates the incoming address format into a form suitable for the fast operation of the core switch and segments the packet to a length suited to the core switch. The core switch consists of three main parts; an input alignment part, an optical interconnect and an electrical fan-out. The optical interconnect may involve optical fan-out or a one-to-one array of connections. The output packet handler reconstitutes the packet into its original (incoming) format. The efficient operation of the electrical fan-in part is achieved by phasing the start of the incoming packets and the electrical fan-in part contains the storage means to provide the necessary queuing for outgoing packets.

Features of interest are the architecture as a whole, the method of phasing packets, the organisation of the electrical fan-in circuits, and the use of Sectorised Orthogonal transfer (SOT).

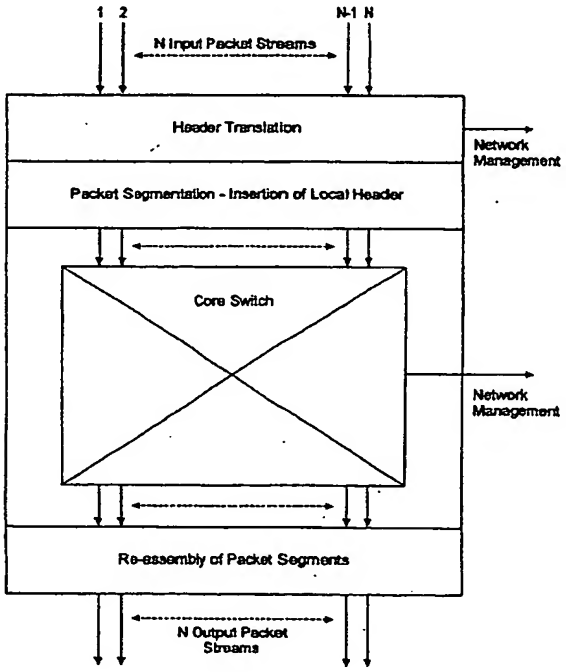

Figure 7 - Packet Router

GB 2 357 656 A

AM

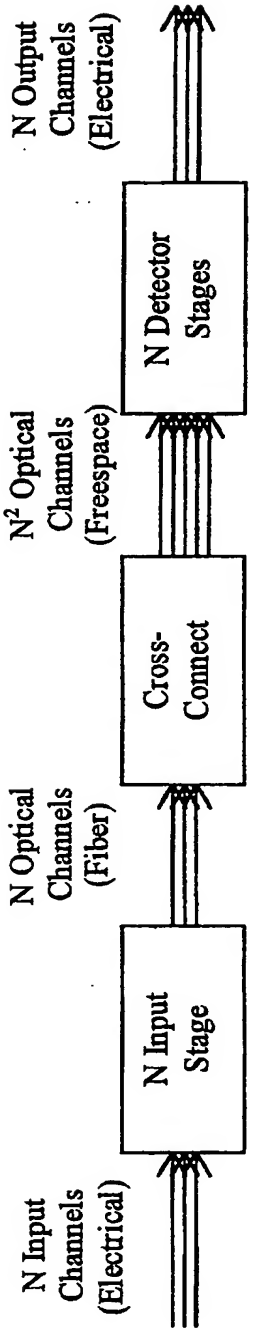

Figure 1. Basic Architecture

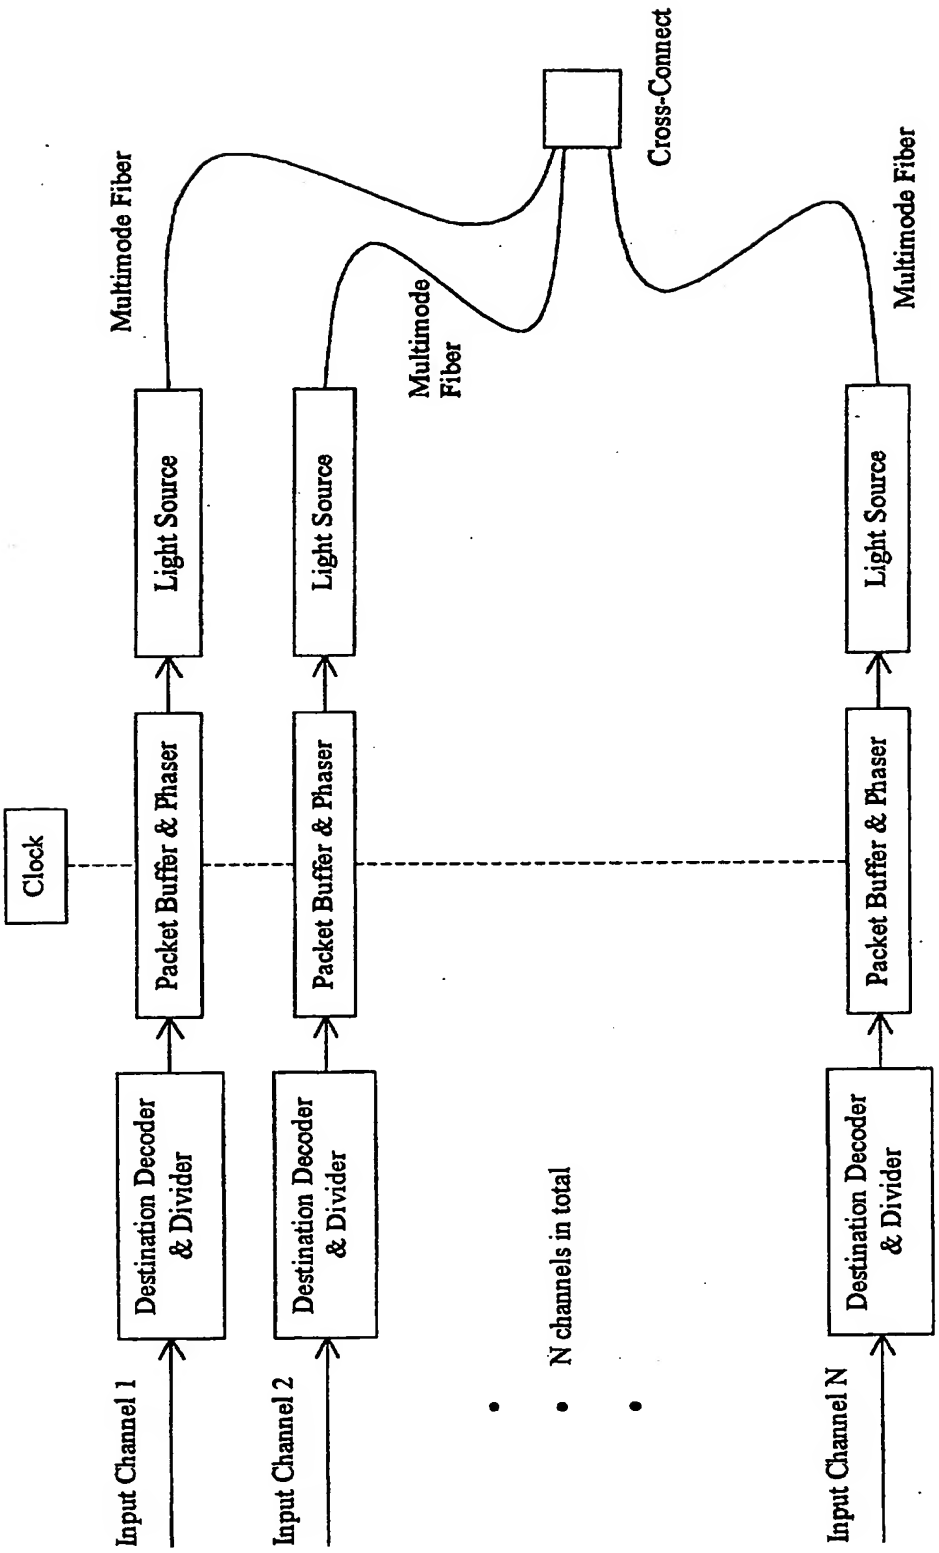

Figure 2. N Input stages

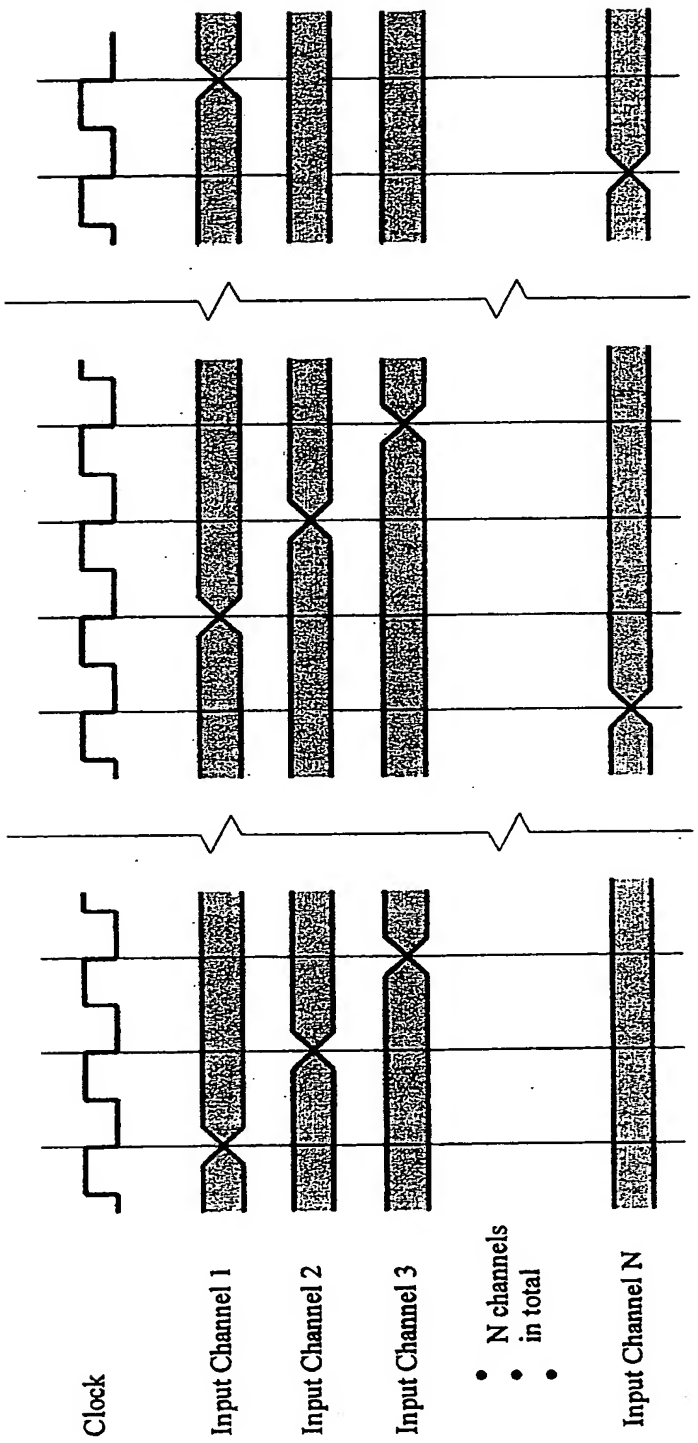

Figure 3. Phasing diagram for serial packet transmission through cross-connect

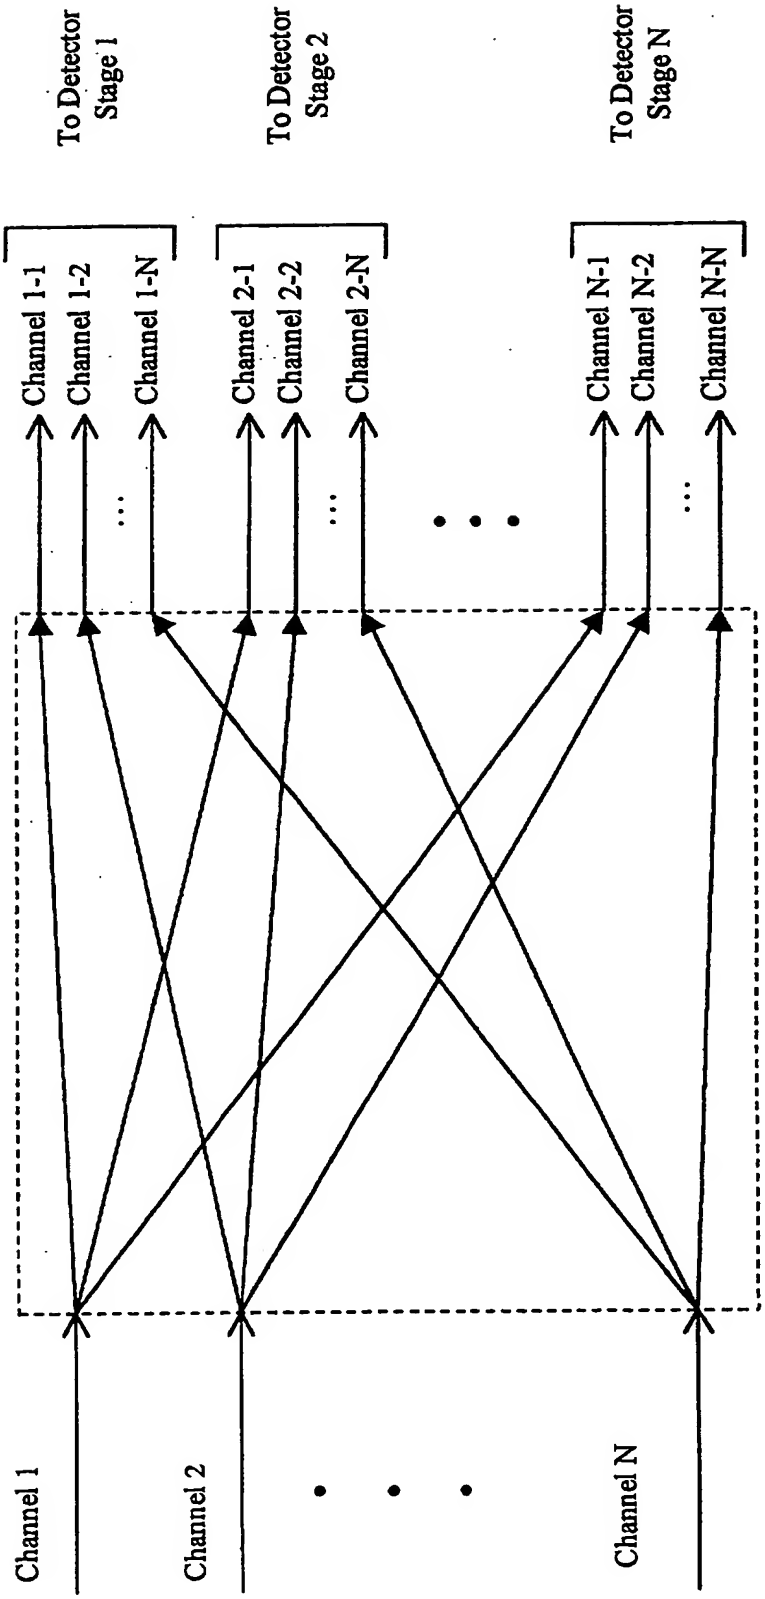

Figure 4. Cross-Connect

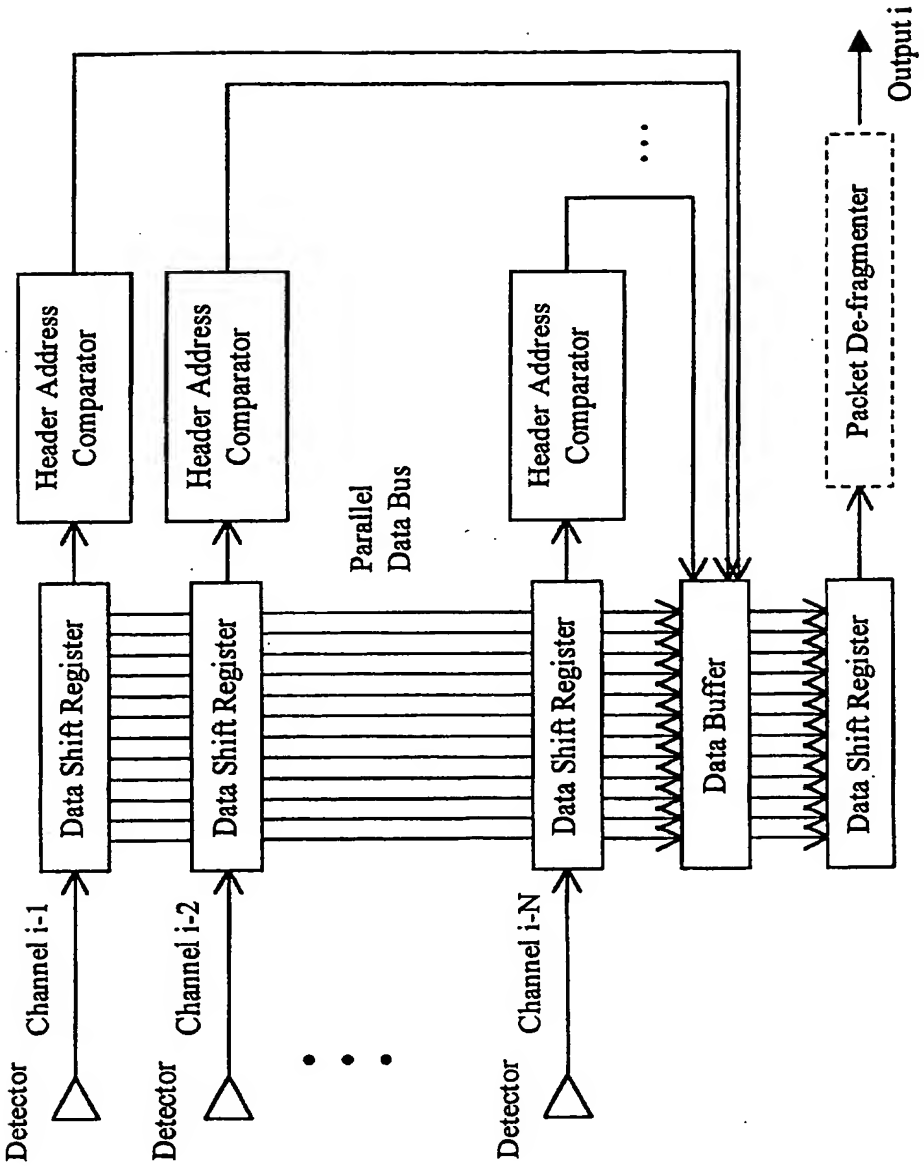

Figure 5. Detector Stage i (N required)

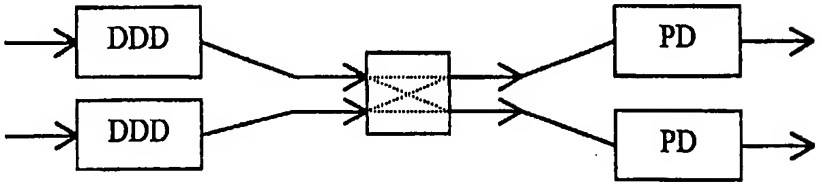

Basic 2×2 router

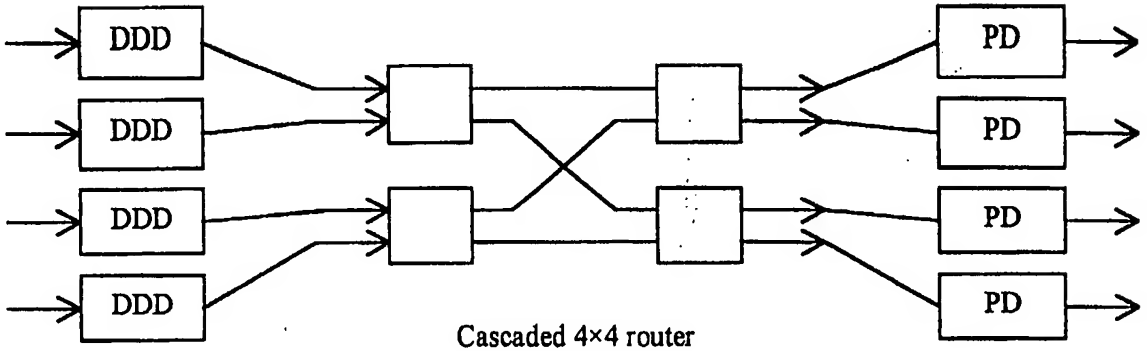

Cascaded 4×4 router

Figure 6. Cascading principle

DDD – Destination Decoder & Divider
PD – Packet Defragmenter

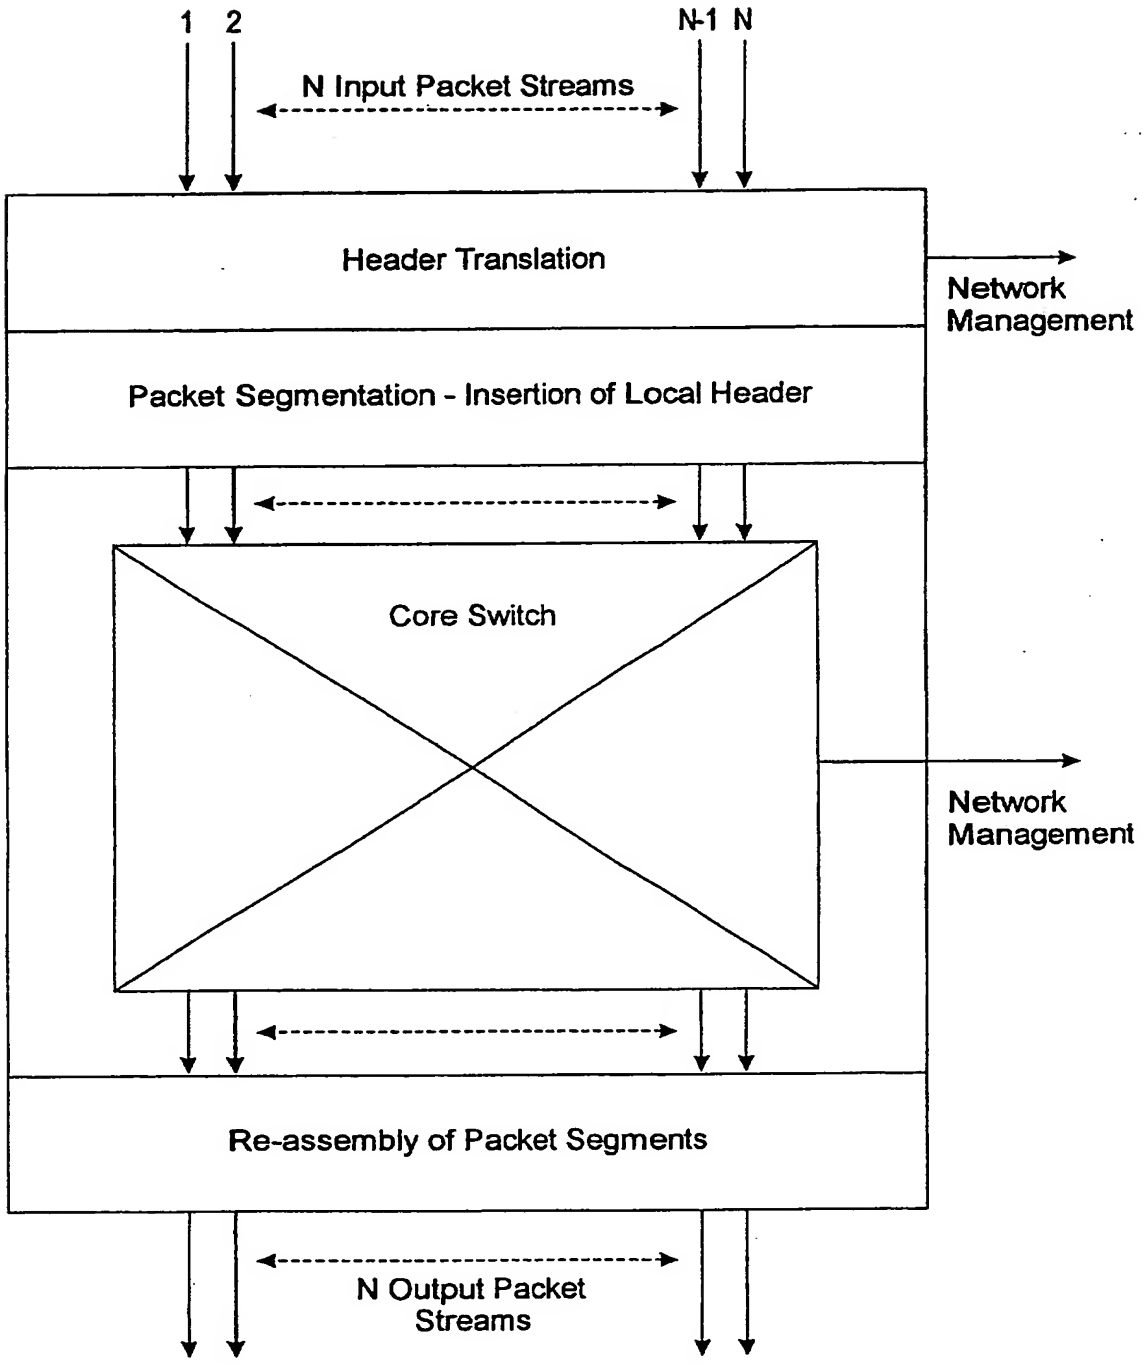

Figure 7 - Packet Router

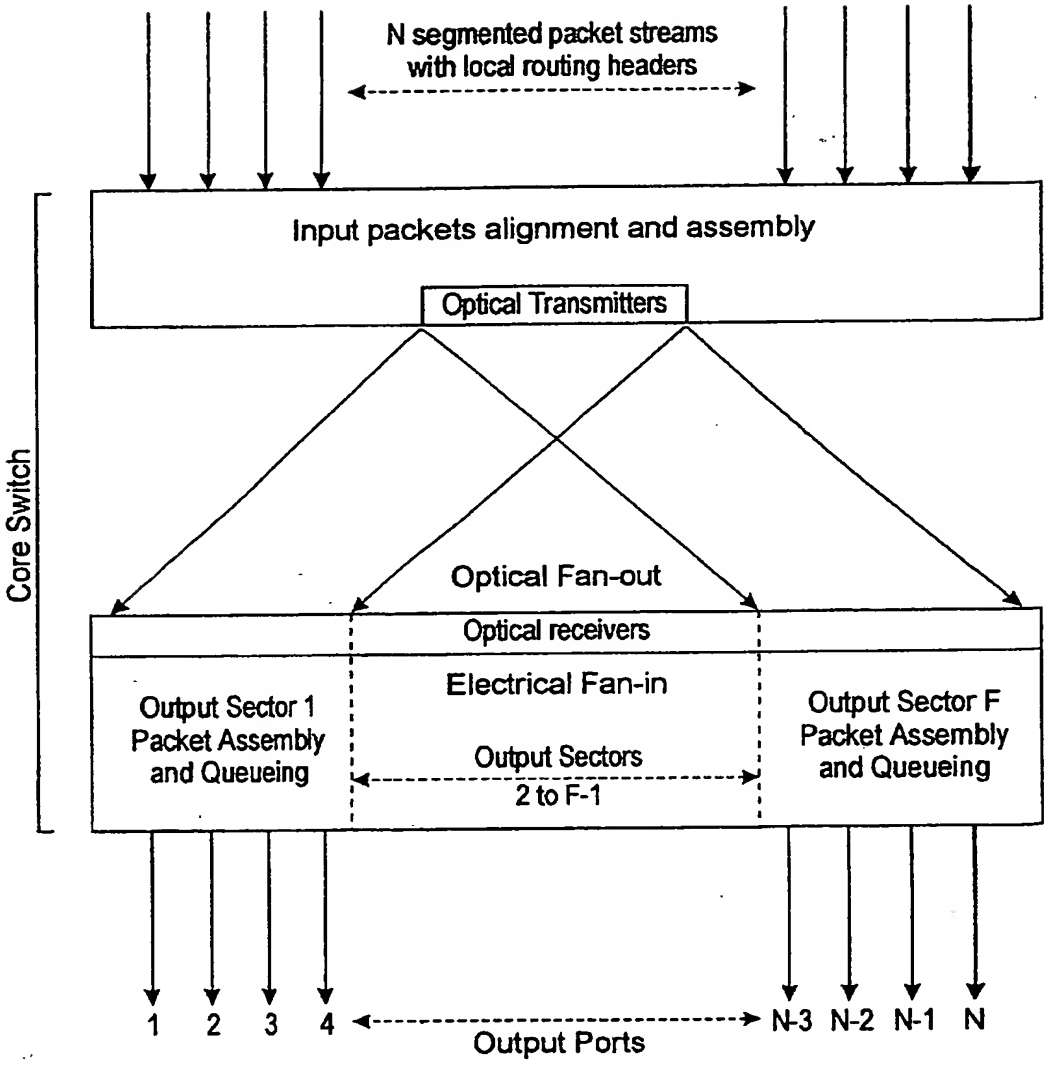

Figure 8 - Block Schematic of the Fan-out Version of the Core Switch

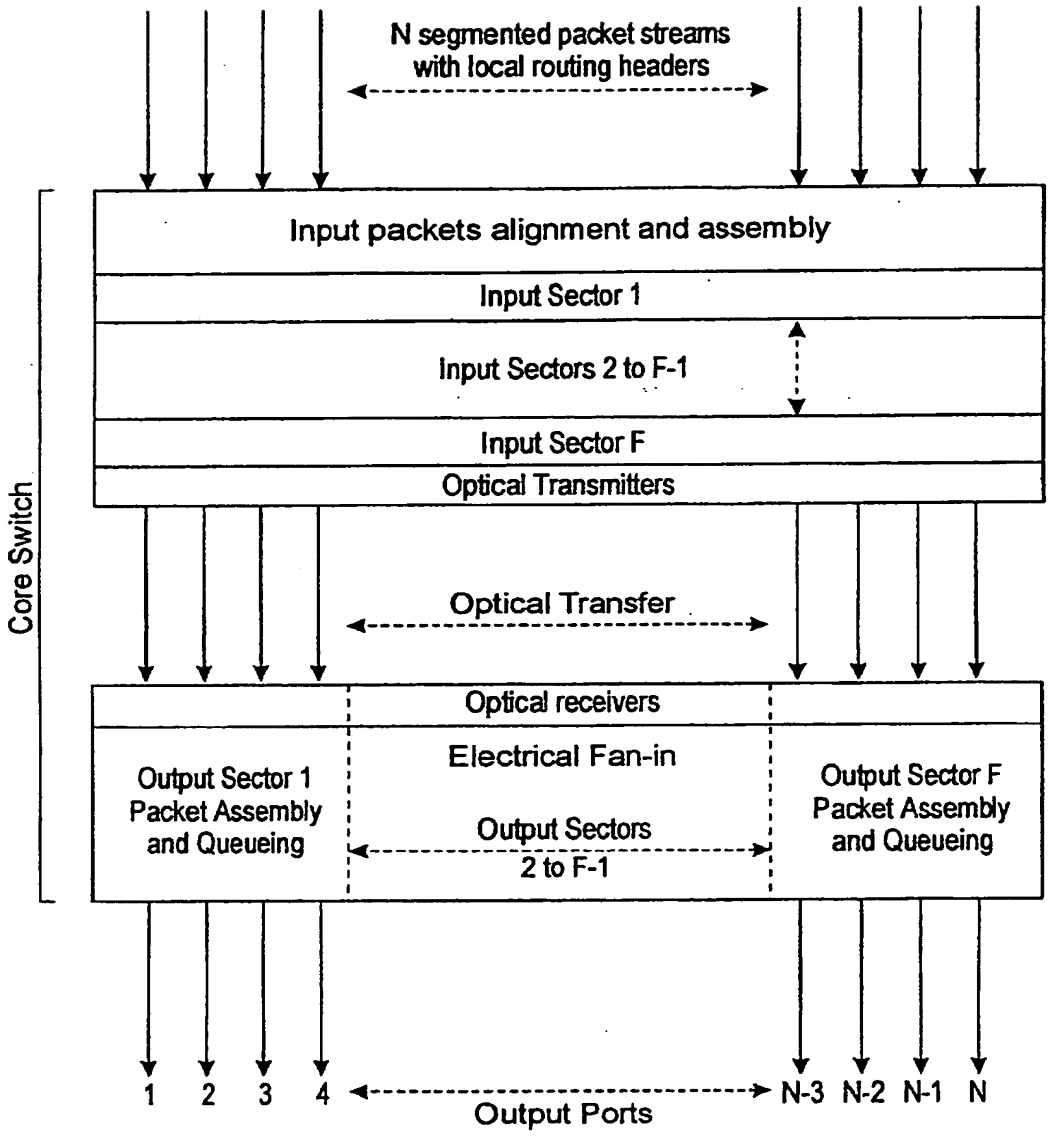

Figure 9 - Block Schematic of the Sectorised Orthogonal Transfer Version of the Core Switch

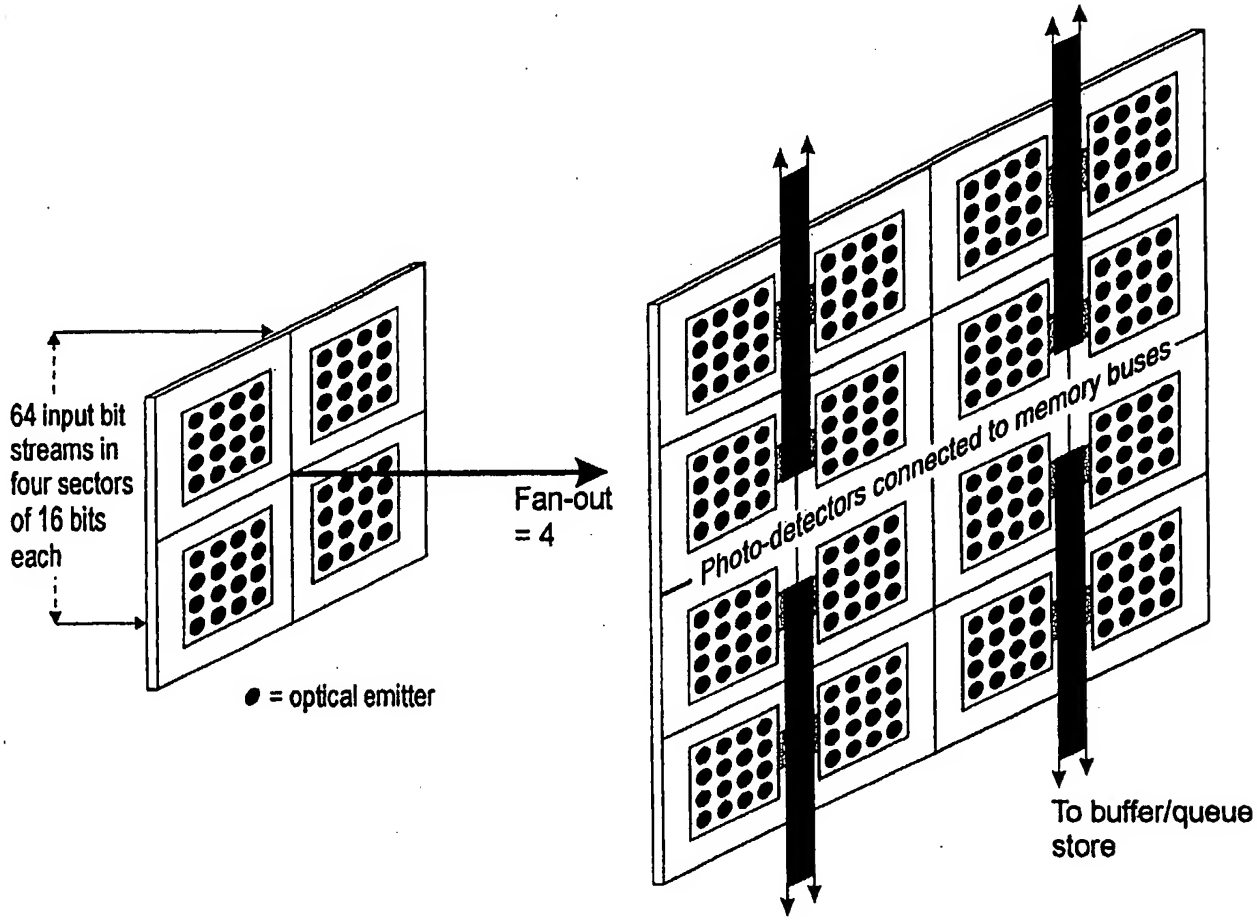

Figure 10 - Principle of Straight Optical Fan-Out (SOFO) with Electrical Fan-in

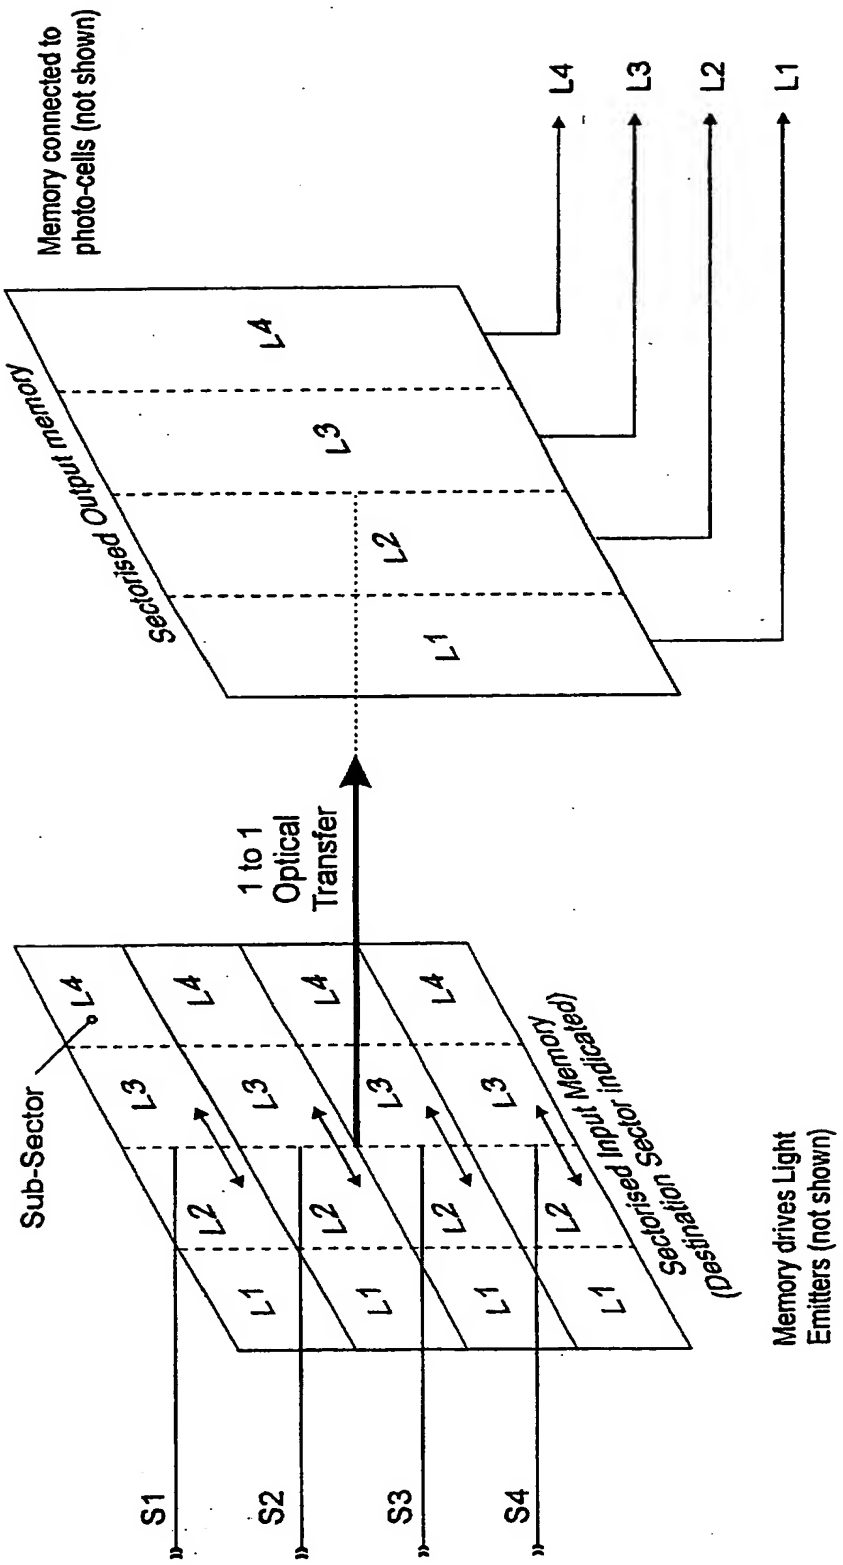

Figure 11 - Sectorised Orthogonal Transfer (SOT)

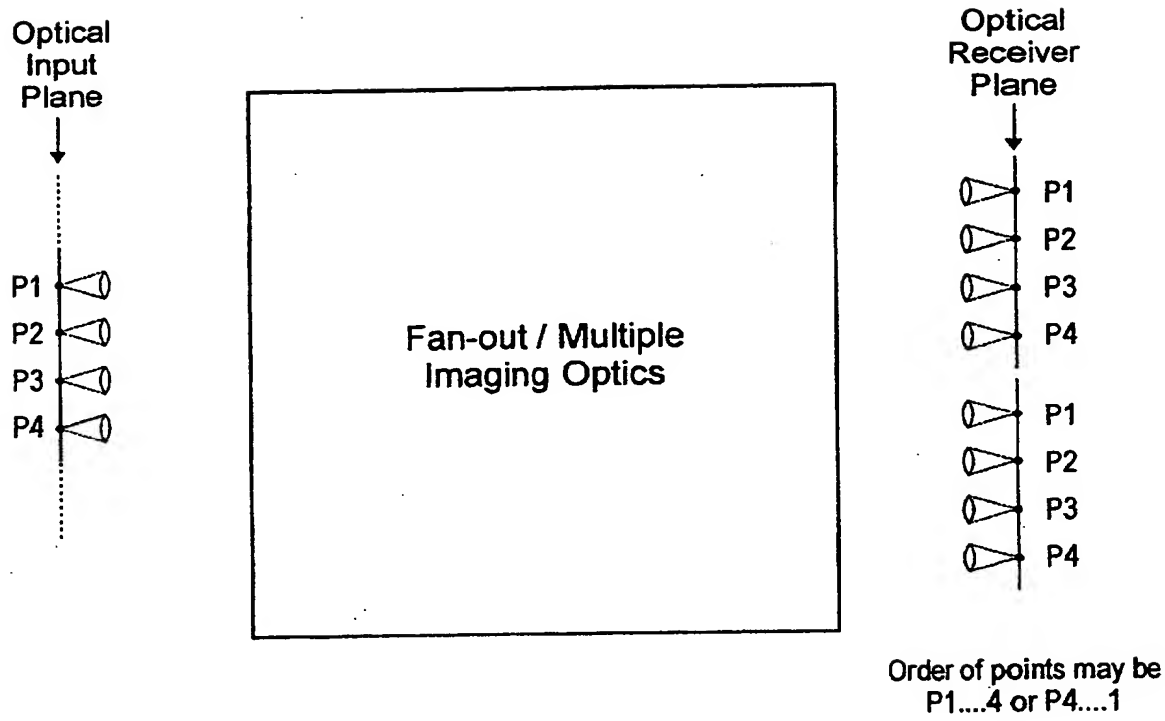

Figure 12 - Schematic Overview of Optical Fan-out / Multiple Image Function

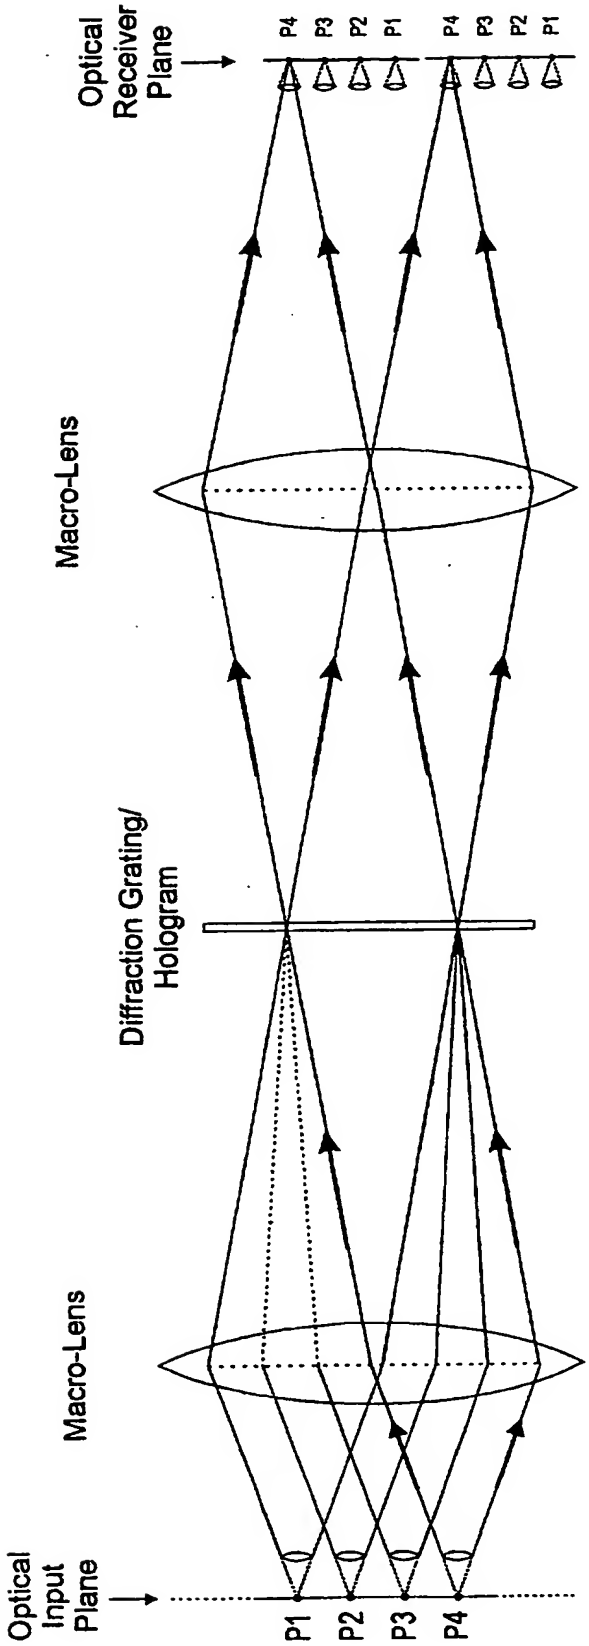

Figure 13 - Multiple Imaging with a Diffraction Grating providing fan-out

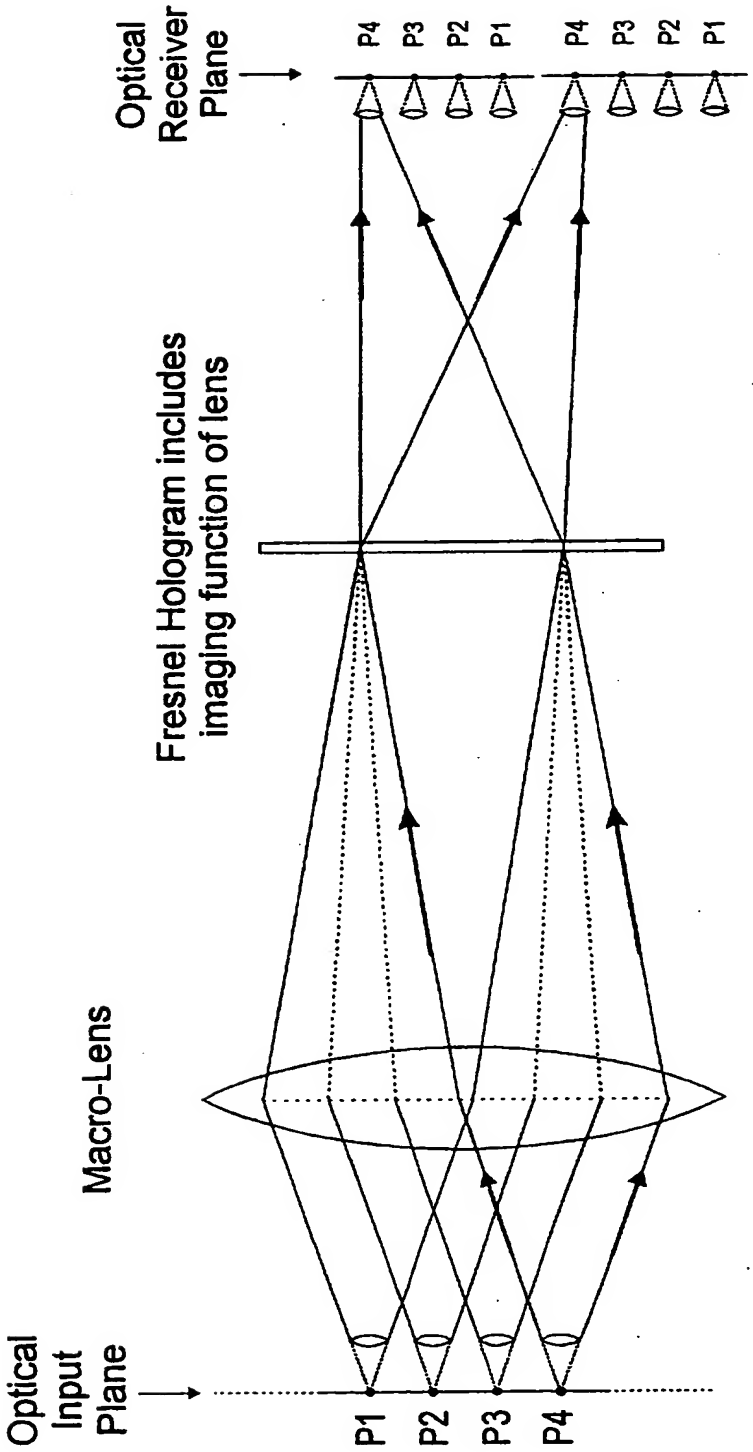

Figure 14 - Multiple Imaging with Fresnel Hologram based fan-out

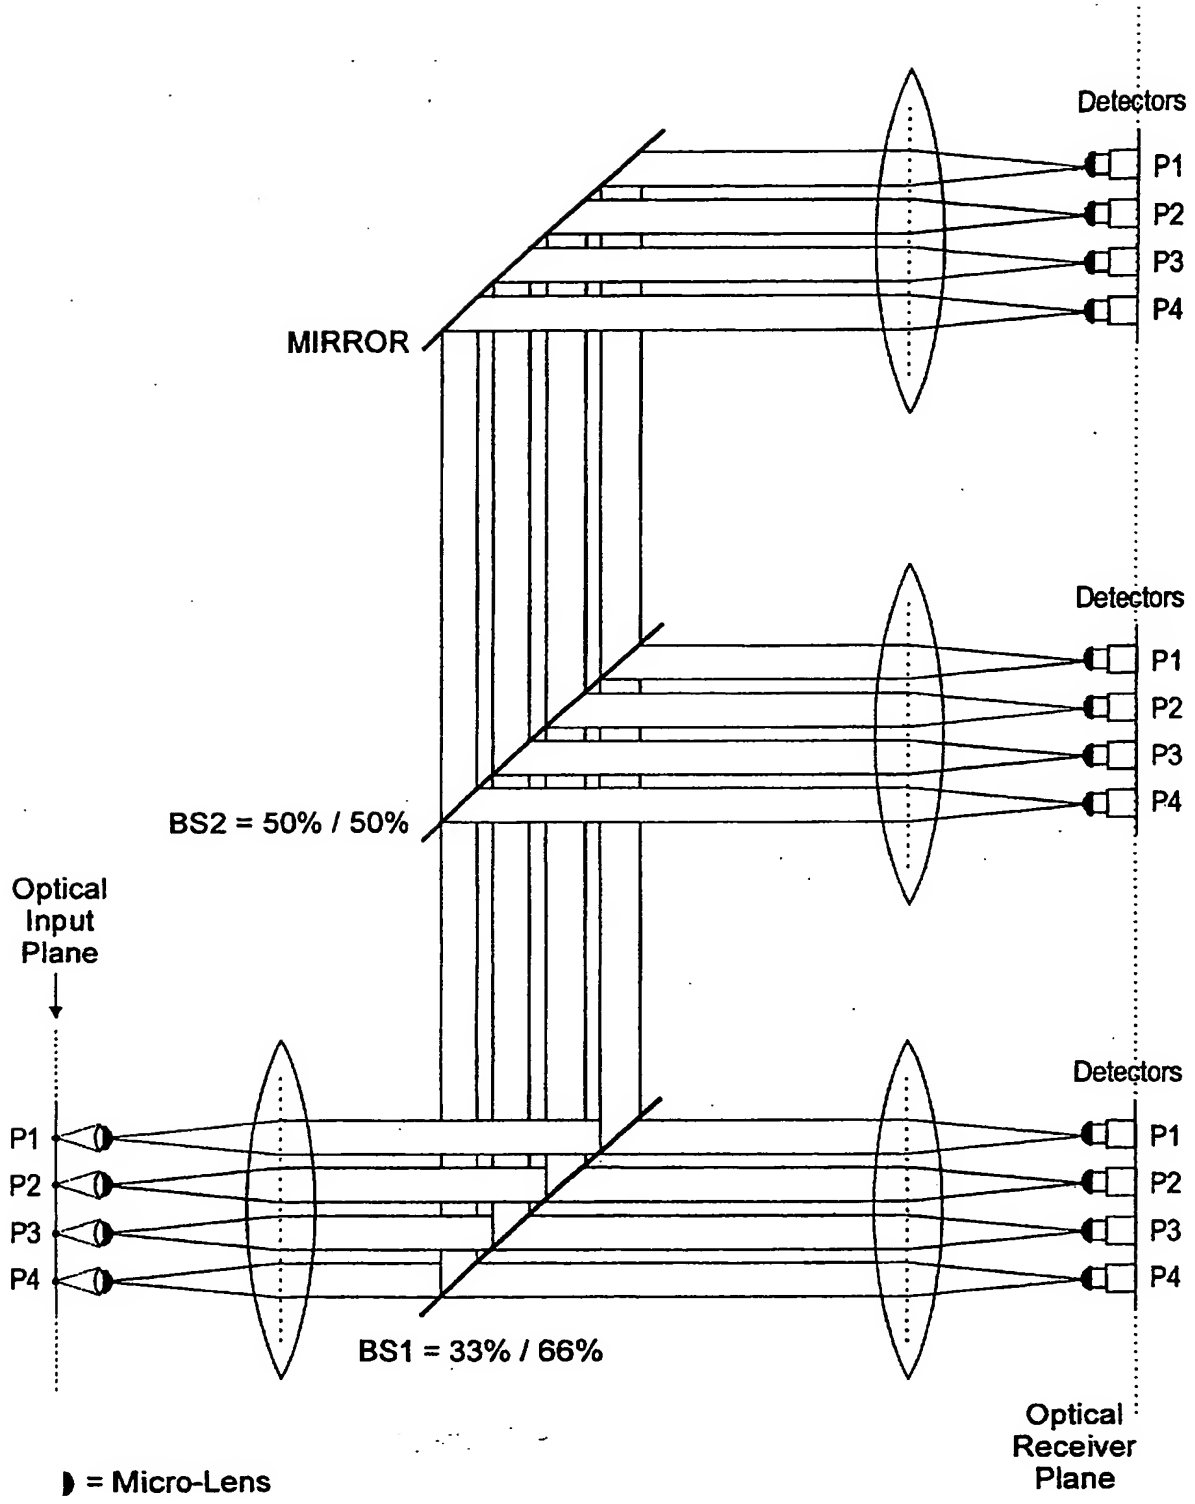

Figure 15 - Beam Splitter based Optical Fan-out / Multiple Imaging

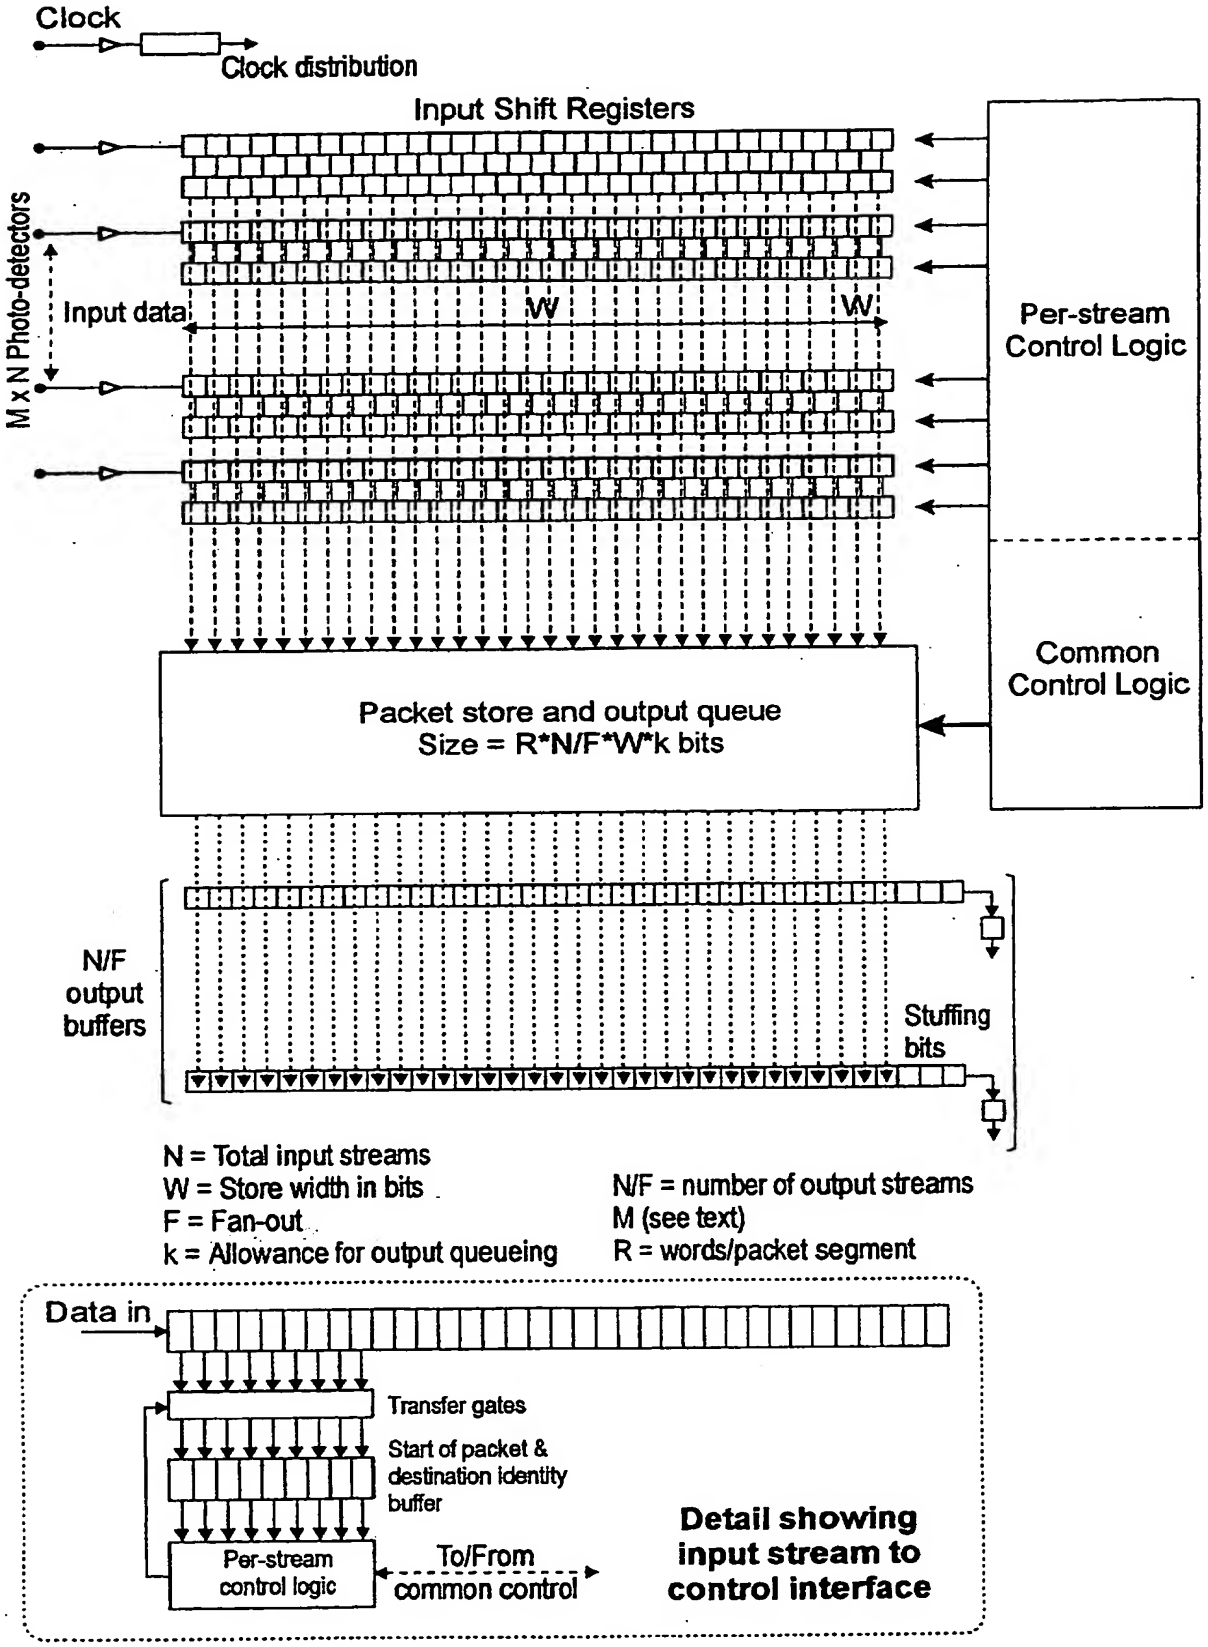

Figure 16 - Block schematic of per sector electrical fan-in circuitry

Router

This invention is concerned with routers, especially those for use in digital communication networks.

Digital communication networks, for example the internet, consist of a web of data channels joined by "routers". These routers switch incoming data packets to their destination according to the "address header" attached to the data packet. Switching many packets between a large number of channels at high speed is difficult and expensive. This invention relates to a design for a packet switching router that enables higher throughput of data at lower cost than existing technology. Moreover, the invention allows routers to be scaled to larger sizes than existing technology.

In the past, the need for faster routers has been tackled by purely electronic means. However, this makes the router difficult to scale, since the cross-connect of N inputs to N possible outputs requires an N^2 matrix of switches which must be addressed. The invention deals with the problem in quite a different way - a "mechanical", and specifically opto-electronic, way. More specifically, the router's required ability to cross-connect is achieved by converting the N electrical signals to light signals. These light signals are "fanned out" so that each is distributed to N detectors, and phased so that only one of the N optical signals is received by the N detectors at any one time. Circuitry associated with each detector reads the header and retains or discards the packet accordingly. The retained packets are then sent through the output channel. The main advantage of free-space optical distribution of the signals it reduces the electrical connection complexity drastically.

the use of an optical cross connect to provide a simple, easily aligned mechanism for the free-space (or other bulk medium) distribution of N signals onto a plane of N^2 photodetectors on an

integrated circuit, while the use of phasing packets through the cross-connect to ensure no packet or sub-packet completes its arrival at the same time as any other in a particular output stage. This ensures there is no contention between incoming packets.

In one aspect, therefore, the invention provides an opto-electronic router including input and output packet handlers that respectively convert between electronic and optical signals and *vice versa*, wherein:

a "core" switch operatively links the handlers, which switch contains distribution means that in operation supply *all* of the input packet handler's optical signal outputs to *all* of the output packet handler's optical signal inputs, and

which switch also contains delay means applicable to each optical output such that all the outputs are separated one from another in time;

and wherein

at each output handler there is associated checking means such that only the optical signals intended for that output handler are converted back to electronic signals and sent on their way.

In a related aspect, the invention provides a method of routing packetised electronic signals between a plurality of input and output lines, in which method the packets from the input lines are converted to optical signals which are then distributed to the output lines where they are back-converted to electronic signal packets for subsequent transmission, which method includes the steps of;

converting the electronic packets from each input line into an optical signal, this signal including an address defining the intended output line;

supplying this optical signal to *all* of the output lines simultaneously. at each of which it is back-converted to an

electronic signal, the optical signals from all the input lines being so supplied sequentially;

and thereafter at each output line checking the address of the received signal and passing on only those signals intended for that output line.

The main principles of the invention are now described for convenience with reference to the accompanying diagrammatic Drawings in which:

Figure 1 shows the basic architecture of the router arrangement;

Figure 2 shows the router's input stages;

Figure 3 shows a phasing diagram for serial transmission through the router;

Figure 4 shows the concept of "fanning out" the signals, from N optical outputs to all of the N optical inputs across the core switch;

Figure 5 shows conceptual details of the detector stages; and

Figure 6 shows the idea of cascading routers of the invention.

The overall architecture is shown in Figure 1. The inputs pass through input stages which identify each incoming packet's destination and convert the packet to a light signal. These light signals are then pass through a cross-connect device which fans out each input channel into N channels (N^2 channels in total). These channels are then taken to a detector stage were the signals are electrically detected and either discarded or retained according to whether they reached the desired destination. The retained signals are buffered and sent to the channel's output.

Each input stage (Figure 2) reads each incoming internet packet's destination header and then decides to which output channel the packet should be routed. (This technology exists in the art.) This information is appended on the packet (the "local header"). If the packets are of uneven length, or shorter packets are optimal, they may be divided into sub-packets at this stage. Information about such fragmentation of the packet will be included in the local header.

The signals in each channel are now buffered and delayed for a short period to ensure that packets or sub-packet in each channel will never complete their arrival at the output stage synchronously (Figure 3 shows an example of a suitable timing diagram for packet transmission). This it to avoid contention at the output stage.

The packets or sub-packets are then converted to serial optical signals, e.g. by using Vertical Cavity Semiconductor Lasers (VCSELs) or light emitting diodes (LEDs) and are connected to the cross-connect stage using multimode optical fibres. This allows the input stages to be located some distance from the cross-connect and does not require careful alignment during construction.

The cross-connect (Figure 4) fans out each input channels into N separate channels. These N^2 channels are then shuffled so that they are arranged into N output groups of N channels, with one of each input channel in each output group. Various optical fan-out technologies are known in the art, including diffraction-grating holograms (Figure 7), Fresnel holograms (Figure 8) and beam splitters (Figure 9) and fibre splitters (see G.J.Grimes, *et al.* "Polymer-Waveguide-Based Photonic Switch", Optical Engineering, Vol 31, No 10, October 1992, pp 2218-2224). The key function of the cross-connect is to provide a simple, easily aligned mechanism for the free-space (or other bulk medium) distribution of N signals onto a plane of N^2 photodetectors on an integrated circuit.

The photodetectors form the first part of the detector stage (Figure 5). Each group of N signals, one from each input channel, is processed by one detector stage. Once each signal has been detected, the data are loaded into a shift register. Since a signal from any input channel is sent indiscriminately to all detector stages, the signal must be discarded by all but one of the detector stages (this is known as the knockout principle - see M.J. Karol, M.G. Hluchyj. 'The knockout switch. Principles and performance.' Proc. 12th Conference on Local Computing Networks, October 1987, pp 16-22). This is achieved by reading the "local header" appended by the input stage. If the signal has reached its correct destination, the local header is discarded and the signal is sent via a parallel data bus to a data buffer. The data buffer queues the signal for re-transmission as a new serial data packet. By implementing the entire detector stage (and possibly all detector stages) on single integrated circuit, interconnection is reduced to a minimum.

If the packets have been divided into sub-packets, they must be recombined. This may be achieved as a separate packet defragmenting component, or in the data buffer and/or data shift registers.

The design of the router lends itself to cascading in order to achieve greater interconnectivity. The principle is shown in Figure 6, where four N=2 routers are cascaded to produce an N=4 router. Note that the first component of the system, the destination decoder (in Figure 2), is required only at the beginning of the entire system and need not be replicated at every stage. (Note that the local header must then not be discarded until the final stage.) Similarly, the packet defragmenter (shown dotted in Figure 5) is required only at the final output stage.

The particularly important elements of the invention are believed to be:

the possible use of VCSEL laser to generate the optical signals;

the use of multimode optical fibres to collect light from remote sources and transport it to the input plane of the cross-connect; and especially

the ability to cascade devices to achieve greater interconnectivity (this provides more channels per unit equipment and higher throughput bandwidth).

In summary, the invention discloses a packet routing architecture consisting of an input packet handler, a 'core' switch and an output packet handler. The input packet handler translates the incoming address format into a form suitable for the fast operation of the core switch and segments the packet to a length suited to the core switch. The core switch consists of three main parts; an input alignment part, an optical interconnect and an electrical fan-out. The optical interconnect may involve optical fan-out or a one-to-one array of connections. The output packet handler reconstitutes the packet into its original (incoming) format. The efficient operation of the electrical fan-in part is achieved by phasing the start of the incoming packets and the electrical fan-in part contains the storage means to provide the necessary queuing for outgoing packets.

The novel and inventive features of this are the architecture as a whole, the method of phasing packets, the organisation of the electrical fan-in circuits, and the use of Sectorised Orthogonal transfer (SOT).

There is an increasing demand for 'routers' in packet switching networks carrying internet traffic. The invention proposes to achieve routing using switching modules of novel design and having a substantial but limited capacity that are easily connected together to form larger entities capable of routing traffic up to many hundreds, if not thousands, of g/bits.

The invention in its apparatus and method aspects may alternatively be defined as follows:-

As regards the method, the invention is a method of routing data in a packetised form between a plurality of incoming links and a plurality of outgoing links, including the conversion of the address data in the packet headers into a reduced form suited to the fast connection of packets by a core switch whereby all incoming data is available for a selected output link, including the segmentation of packets into sub-packets of a length suited to the fast transfer of packets by the core switch, the said core switch including means to phase the packets such that the start of a sub-packet is allocated a specific time slot in a multiplex and means to add the reduced form of address to the sub-packet to facilitate its routing through the core switch, means for the optical transfer of the said sub packet in serial

form to optical detectors coupled to an electronic output unit, said electronic output unit containing means for converting the sub packets from serial form to parallel form and transferring them to a memory store, means for holding sub packets in the memory store until the outgoing link for which they are destined is free of traffic, means for reading out the memory store and transforming the sub packets back to serial form, means for the removal of the reduced header and the reassembly of complete packets before transmission on an outgoing link.

As regards the router architecture - the apparatus - the invention is defined as the hardware inherent in the defined method, in which hardware there are N replications of the electronic output units each unit being a called a sector and providing output to only 1/N th of the output links, optical transfer means providing an N-fold replication of the complete input data to each of the N sectors each of which said sectors ignores the data not wanted by it. In a preferred form, there are N replications of an electronic input unit and N replications of an electronic output unit each input and output unit being divided into N sub sectors, optical means for orthogonally connecting sub sectors of input units to sub sectors of output units such that a sub sector of an

input unit with coordinates a,b is connected to an output unit sub sector b,a (where the first coordinate refers to the sector and the second to the subsector), means for discarding at each input unit sub sector packets which are not intended for the output sub sector to which it is connected.

And, of course, the invention allows cascading, providing a packet switching architecture in which a multiplicity of architectures as in claims 1,2 and 3 are interconnected as modules to make a larger packet switching architecture, means for carrying the reduced form of addressing and sub packeting over several architectures as defined just above in cascade.

Embodiments of the invention are now described, though by way of illustration only, with reference to the accompanying diagrammatic Drawings in which:

Figure 7 shows a packet router of the invention;

Figure 8 shows a block schematic of the fan-out version of the core switch of the router of Figure 7;

Figure 9 shows a block schematic of the Sectorised Orthogonal Transfer (SOT) version of the router of Figure 7;

Figure 10 shows the principle of the use of straight optical fan-out (SOFO) with electrical fan-in;

Figure 11 shows the principle of the use of the alternative sectorised orthogonal transfer (SOT);

Figure 12 shows a schematic view of optical fan-out, with multiple image function;

Figure 13 fan-out and multiple imaging using a diffraction grating;

Figure 14 shows optical fan-out based on a Fresnel hologram;

Figure 15 shows optical fan-out effected using a beam-splitter; and

Figure 16 shows a schematic of per sector electrical fan-in circuitry.

Figure 7 shows a block schematic of the module design. It consists of an input packer handler, a core switch and output packet handler. The input packet handlers perform two functions:

(a) receive packets routed according to any one of a number of standard protocols (e.g. IP (Internet Protocol), Ethernet, ATM or frame relay). Associated with these packet handlers are look-up tables which translate the packet header information into a local address in the form of a header. This address contains fields that identify the route to be

taken though the core switch. Compared with the packet header, these address fields are very short, making for fast route determination.

(b) The packet is segmented into a number of octets, PS, which are chosen for optimal core switch size/performance. The value of PS will be much less than the maximum IP packet length and might be of the same order as an ATM packet (PS = 53 octets).

The Core Switch

There are number of possible architectures for the core switch, only two of which are described here. They all have the same overall functionality but differ in the detail of the design. The switch is based on a 'broadcast' or 'the generalised knockout principle. Figure 8 shows one of the embodiments based on straight optical fan-out (SOFO) and Figure 3 shows the other based on sectorised orthogonal transfer (SOT).

The principle of Figure 8 is amplified in Figure 10 for a particular example of 64 inputs and an optical fan-out of 4. In general the inputs are copied F times to F output sectors. F is the fan-out which is best achieved optically avoiding the cross coupling problems associated with an all electronic solution. Each of the F sectors receives the complete set of input packets but ignores/discards those not intended for it. Those that are wanted are written to a per sector store which stores packet segments and is organised to act as a queue that holds those packet segments until their destination ports are free. For example with N = 64 (Figure 4), a fan-out of 4 with 4 sectors each having 16 output ports, the electronics are then performing a fan-in of 64 to 16 i.e. a fan-in of 4.

With sectorised orthogonal transfer (SOT), Figure 9 and as amplified in Figure 11, the input is first divided into sectors and then sub-sectors. The input for one sector is broadcast to all sub-sectors but is only selected by a given sub-sector if that sub-sector is 'opposite' the relevant output sub-sector. Sub-sectors are arranged in an orthogonal array and there is a direct optical connection between sub-sectors on the input and output sides. SOT needs less optical

connections than SOFO but is more sensitive to traffic patterns and hence may lose more packets.

Input Packet Alignment and Assembly

If the incoming data rate per packet is g bit/s then the incoming clock interval is 1/g = ti secs. Each of the N incoming packets is allocated a time slot modulo N and of duration ti - this is done to avoid store contention problems on fan-in. Packets are transferred in parallel but the incoming packet transfer is only allowed to start transfer when its time slot occurs. Hence the input packet assembly buffers must be at least N bits per packet to achieve the desired phasing (see Figure 3). Further segmentation into sub-segments of T = 8*PS(1+OH)/W bits takes place at this point, for the definition of W see the section on electrical fan-in below and OH represents the local addressing over head plus stuffing bits. T is called the 'tranche size'.The local address is added to the first of the sub-segments and one or more stuffing bits are added between each sub-segment. If a packet is present on a given input, its first sub-segment will start transmission in its time slot in serial form from the optical transmitter to the input detectors of the fan-in block.

Optical Interconnects

There are a number of existing technologies for providing the optical interconnection and this section describes some of them; they form part of an overall architecture which is believed to be novel.

Figure 12 provides a non-technology dependent view of the function of an optical fan-out architecture. The plane labelled Optical Input Plane contains point sources organised as a one or two dimensional array and could be either;

(a) An image of spots transported to that plane from non-local emitters via an optical system which reduces the spatial lateral distance between the individual emitters.

(b) A number of spatially distributed optical emitters, which could be organised as a one or two dimensional array.

The SOFO design fits into the optical fan-out section of the core switch (Figure 8) and requires that the array of point sources in the Optical Input Plane (OIP) are copied to multiple copies of the array displaced spatially and non-overlapping at the Optical Receiver Plane (ORP), as labelled in Figure 13. The construction of each copied array may be magnified positively or negatively in comparison to the OIP in order that the size of the point sources matches the size of the individual detectors in the ORP. It is possible that the copied images may be copied without maintaining the same orientation as the process of fan-out may cause each copy to be "flipped" round its local axis. This would result in points PI..4 being ordered as P4.. 1, etc. This factor is an artefact of the optical system design and results as a consequence of different optical paths for a split image. It is not claimed that the points within an individual copy are re-ordered amongst themselves other than through this flipping process, thus, a point in the array will maintain its neighbourhood.

There are a number of processes by which the light from each source in the array can be split into multiple copies of itself, including free-space optical components such as beam splitters, diffraction gratings, thin and thick holographic components. For low orders of fan-out F (less than eight in one-dimensional fan-out and less that sixty four in two dimensional fan-out), it may be feasible to use waveguide light splitting devices. The latter includes passive waveguide structures that split the light propagating along the input waveguide into more than one output waveguide such as those constructed in integrated passive waveguide or optical fibre technologies. A fibre splitter implemented by Grimes *et al* allows the transfer of light from a large core diameter (1 mm) multi-mode plastic fibre to a bundle of fibres with a 200um diameter core. A Perfluorinated Ethylene Propylene sheath is used to fix the multi-mode fibre and the fibre bundle in position.

The optical design of the waveguide based optical fan-out system would require additional optics in order to control the beam divergence

and spot size of each of the point sources propagating from the OIP. The optics could be as simple as a single microlens designed to couple the light from one point source into one waveguide input channel. The waveguide would then confine the light to propagate along its channel and split according to the construction and placement of the waveguide splitter technology.

In the case of the holographic free space systems, it is adequate to image the points using a common-to-all macrolens such that the resolvable points in the OIP plane are imaged to an overlapping point which illuminates a sufficient extent of the diffractive/holographic component (Figure 14). This component is designed to control the splitting of the light into the multiple images and may require a lens to be used to image the multiple images to the ORP (as shown in Figure 14). Alternatively, the diffractive/holographic component may itself incorporate an imaging function and image the light to the ORP (as shown in Figure 15),

In the case of a beam splitter (BS) based fan-out system, each BS will split the light of the array into two identical arrays, in terms of their spatial/angular light distributions, but not necessarily in terms of the optical power split. A cascade of BS components (which may not be identical in their power splitting behaviour) will define the optical paths between the OIP and the ORP. Each copy of the OIP will follow one particular optical path through the BS network. As an example, consider Figure 15 where the sources at the OIP are imaged through a lens and travel as parallel beams through the BS network. The initial beam splitter (BS 1) splits the light into two paths such that 33% of the optical power travels straight through the beam splitter and out towards the OR-P. The remaining 66% of the power is deflected to upwards and into BS2 which splits this light into two paths with 50% of the power travelling in each direction. One path leads directly to the ORP whilst the other first reflects from a mirror, which reflects all of the light on towards the OR-p. The optical design shown here uses microlenses to collimate the light from the points in the OIP and it uses macrolenses to image the beams through the beam splitter network. This example

illustrates that BS can be used to perform multiple imaging and that, if desired, the optical power splitting performed by the beam splatters can be designed to allow an approximately even sharing of the optical power between copies of the input image,

It is not necessarily the aim of the multiple imaging optics to split the optical energy of a source in the OIP array evenly between its fanned-out points at the ORP.

In all cases, the process used to control the multiple imaging of the OIP array to the ORP should aim to communicate the optical energy from a point source in the OIP to its corresponding receivers in the output plane. Any optical energy which is imaged to the OIP which does not arrive at its correct receiver area will be seen as optical crosstalk between channels in the switch. The spot size of each of the fan-out spots must be controlled by the optical system in order that the major proportion of its light will impinge on a single optical detector in the ORP.

Electrical Fan-in

The general organisation of the electrical fan-in is shown in Figure 16 and is described in the context of SOFO. (The operation for SOT is very similar except that the unwanted packet segments have been previously discarded.) For a specific sector the packet segments are received by the optical detectors (Figure 8) and are loaded serially into shift registers and the local header examined. If 'not for me' the packet is ignored and drops off the end of the shift register. The shift register length is W bits which is also the width of the sector queue store. When W bits are received these are transferred in parallel to the queue store. To cope with N incoming time slots, the sector store write time needs to be W/Ng secs. Hence W can be chosen to match the speed capability of the storage technology. The purpose of the inter sub-segment stuffing bits is to allow the serial shifting to stop without loss of information while parallel read-out takes place from a shift register into the queue store. Part of the local address is the

destination port number of the output and at each output port there is a queue of pointers. At the start of a new packet, the address of the queue store location assigned to it is put on the relevant pointer queue. The number of locations in the queue store per packet is T, the tranche size. Segment size, tranche size and W can be adjusted within limits to match the capability of the core switch - in particular the speed/size performance of the queue store can be optimised. Each output stream is assigned a time slot modulo N/F with read time available of N WF/Ng secs. The output process is the converse of the input, i.e. parallel read-out from the queue store with parallel read-in into a shift register, one per port. The stuffing bits are re-inserted and serial read out takes place from the shift registers.

Output

For a stand alone module, output can be into a packet re-assembler, which involves stripping off the stuffing and local address bits and reconstituting the full packet from packet segments in the form appropriate to its outgoing network protocol. When a module forms part of a larger switch and the output of a first stage module is the required input for a similar second stage module, then the local header can be designed to embrace multistage operation and no stripping is necessary at that point.

CLAIMS.

Claim 1. A method of routing data in a packetised form between a plurality of incoming links and a plurality of outgoing links, including the conversion of the address data in the packet headers into a reduced form suited to the fast switching of packets by a core switch whereby all incoming data are available for selected output links and including the segmentation of packets into sub-packets of a length suited to the fast transfer of packets by the core switch the said core switch including means to phase the packets such that the start of a sub-packet is allocated a specific time slot in a multiplex and means to add the reduced form of address to the sub-packet to facilitate its routing through the core switch, means for the optical transfer of the said sub packet in serial form to optical detectors coupled to an electronic output unit, said electronic output unit containing means for converting the sub packets from serial form to parallel form and transferring them to a memory store, means for holding sub packets in the memory store until the outgoing link for which they are destined is free of traffic, means for reading out the memory store and transforming the sub packets back to serial form, means for the removal of the reduced header and the reassembly of complete packets before transmission on an outgoing link.

Claim 2. A packet switching apparatus as in Claim 1 in which there are N input ports, M output ports (N less than equal to or greater than M) and F replications of the electronic output units each unit being a called a sector and providing output to only M/F of the output links, with optical transfer means providing an F-fold replication of the N input data streams to each of the F sectors each of which said sectors discards the data not destined for it, means for storing within said output sectors all the packets destined for the said sectors until such time as the output links to which they are addressed are free to transmit the said packets.

Claim 3. A packet switching apparatus as in claim 1 in which there are N input ports, M output ports (N less than equal to or greater than M) and F replications or sectors of an electronic input unit and F replications or sectors of an electronic output unit each input and output unit being divided into F sub sectors, optical or electrical means for orthogonally connecting sub sectors of input units to sub sectors of output units such that a sub sector of an input unit with coordinates a,b is connected to an output unit sub sector b,a wherein the first coordinate refers to the sector and the second to the sub-sector, means for disregarding at each input unit sub sector packets which are not intended for the output sub sector to which it is connected and means for storing within one output sector all the packets received by its own sub sectors until such time as the output link to which the said packets are addressed is free to transmit the said packets.

Claim 4. A packet switching apparatus as in claim 3 in which the number of connections between input electronic units and output electronic units is greater then or equal to N the number of input ports.

Claim 5. A packet router in which a multiplicity of packet switches as in claims 1,2 and 3 are interconnected as modules to make a larger router, means for carrying the reduced form of addressing and packet segmentation as in Claim 1 over several apparatus as in Claims 1,2,3 and 4 in cascade.

Claim 6. A packet switching apparatus with two blocks of P, say, switches as in claims 2 or 3 and 4 with the first block having P output ports connected to P input ports of the second block wherein the connection means is an optical or electrical orthogonal transpose wherein the connections obey the rule that output port number h say on block j say of the first block of core switches is connected to input port number j on the on block h of the second block wherein h and j are any numbers in the range 1 to P.

Claim 7. Apparatus as in claims 2,3 and 4 in which the means of connection between the electronic input units and electronic output units are optical transmitters in the form of Vertical Cavity Semiconductor Lasers (VCSELs)

connected by any optical means to photo-receivers forming an integral part of a CMOS VLSI circuit.

Claim 8. Apparatus as in claim 3 and 4 in which flip chip bonding is the means of connection between the input electronic units and the output electronic units forming the input and output sectors respectively, means of realising said input and output electronic units singly or in combination in the form of semiconductor chips.

Claim 9. Apparatus using the method of Claim 1 and the packet switching apparatus of Claim 2 in which for a specific sector the photo-receiver's outputs are read serially into shift registers where packets segments not addressed to the said specific sector are discarded, means in the relevant time slot of a multiplex for parallel read out of the content of said shift registers directly or indirectly into a memory store means for holding said packet segments in the said memory store until such time as they are free to be read out to the destination port of the packet switch.

Claim 10. Apparatus using the method of Claim 1 and the apparatus of Claim 3 in which for a specific sub-sector the photo-receiver's outputs are read serially into shift registers, means for combining the outputs of sub-sectors into a combined output for the sector, means in the relevant time slot of a multiplex for parallel read out of the content of said shift registers directly or indirectly into a memory store means for holding segments in the said memory store until such time as they are free to be read out to the destination port of the packet switch.

Claim 11. Switching apparatus substantially as described herein with reference to and as illustrated in the accompanying drawings, in particular figures 2,3,4,5,6,11 and 12.

Amendments to the claims have been filed as follows

Claim 1. A method of routing data in a packetised form between a plurality of incoming links and a plurality of outgoing links, including the conversion of the address data in the packet headers into a reduced form suited to the fast switching of packets by a core switch whereby all incoming data are available for selected output links and including the segmentation of packets into sub-packets of a length suited to the fast transfer of packets by the core switch the said core switch including means to phase the packets such that the start of a sub-packet is allocated a specific time slot in a multiplex and means to add the reduced form of address to the sub-packet to facilitate its routing through the core switch, means for the optical transfer of the said sub packet in serial form to optical detectors coupled to an electronic output unit, said electronic output unit containing means for converting the sub packets from serial form to parallel form and transferring them to a memory store, means for holding sub packets in the memory store until the outgoing link for which they are destined is free of traffic, means for reading out the memory store and transforming the sub packets back to serial form, means for the removal of the reduced header and the reassembly of complete packets before transmission on an outgoing link.

Claim 2. A packet switching apparatus as in Claim 1 in which there are N input ports, M output ports (N less than equal to or greater than M) and F replications of the electronic output units each unit being a called a sector and providing output to only M/F of the output links, with optical transfer means providing an F-fold replication of the N input data streams to each of the F sectors each of which said sectors discards the data not destined for it, means for storing within said output sectors all the packets destined for the said sectors until such time as the output links to which they are addressed are free to transmit the said packets.

Claim 3. A packet switching apparatus as in claim 1 in which there are N input ports, M output ports (N less than equal to or greater than M) and F replications or sectors of an electronic input unit and F replications or sectors of an electronic output unit each input and output unit being divided into F sub sectors, optical or electrical means for orthogonally connecting sub sectors of input units to sub sectors of output units such that a sub sector of an input unit with coordinates a,b is connected to an output unit sub sector b,a wherein the first coordinate refers to the sector and the second to the sub-sector, means for disregarding at each input unit sub sector packets which are not intended for the output sub sector to which it is connected and means for storing within one output sector all the packets received by its own sub sectors until such time as the output link to which the said packets are addressed is free to transmit the said packets.

Claim 4. A packet switching apparatus as in claim 3 in which the number of connections between input electronic units and output electronic units is greater then or equal to N the number of input ports.

Claim 5. A packet router in which a multiplicity of packet switches as in claims 1,2 and 3 are interconnected as modules to make a larger router, means for carrying the reduced form of addressing and packet segmentation as in Claim 1 over several apparatus as in Claims 1,2,3 and 4 in cascade.

Claim 6. A packet switching apparatus with two blocks of P, say, switches as in claims 2 or 3 and 4 with the first block having P output ports connected to P input ports of the second block wherein the connection means is an optical or electrical orthogonal transpose wherein the connections obey the rule that output port number h say on block j say of the first block of core switches is connected to input port number j on the on block h of the second block wherein h and j are any numbers in the range 1 to P.

Claim 7. Apparatus as in claims 2,3 and 4 in which the means of connection between the electronic input units and electronic output units are optical transmitters in the form of Vertical Cavity Semiconductor Lasers (VCSELs)

connected by any optical means to photo-receivers forming an integral part of a CMOS VLSI circuit.

Claim 8. Apparatus as in claim 3 and 4 in which flip chip bonding is the means of connection between the input electronic units and the output electronic units forming the input and output sectors respectively, means of realising said input and output electronic units singly or in combination in the form of semiconductor chips.

Claim 9. Apparatus using the method of Claim 1 and the packet switching apparatus of Claim 2 in which for a specific sector the photo-receiver's outputs are read serially into shift registers where packets segments not addressed to the said specific sector are discarded, means in the relevant time slot of a multiplex for parallel read out of the content of said shift registers directly or indirectly into a memory store means for holding said packet segments in the said memory store until such time as they are free to be read out to the destination port of the packet switch.

Claim 10. Apparatus using the method of Claim 1 and the apparatus of Claim 3 in which for a specific sub-sector the photo-receiver's outputs are read serially into shift registers, means for combining the outputs of sub-sectors into a combined output for the sector, means in the relevant time slot of a multiplex for parallel read out of the content of said shift registers directly or indirectly into a memory store means for holding segments in the said memory store until such time as they are free to be read out to the destination port of the packet switch.

Claim 11. Switching apparatus substantially as described herein with reference to and as illustrated in the accompanying drawings, in particular figures 2,3,4,5,6,11 and 12.

The Patent Office

INVESTOR IN PEOPLE

22

Application No: GB 9916680.3
Claims searched: 1-11

Examiner: B.J.SPEAR
Date of search: 10 April 2001

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.S): H4B(BN);H4K (KTKX)

Int Cl (Ed.7): H04Q 11/00

Other: Online:WPI,EPODOC,JAPIO,INSPEC

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
	NONE	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.